May 5, 2003

DØ VERTICAL LUMINOSITY REVIEW.

These are the findings of the Dzero Vertical Review committee established by the Spokesmen in November of 2002. This document contains the following sections:

Executive Summary

The Hardware of the Luminosity System

The Readout/Simulation of the Luminosity System

The Online Luminosity System.

The Offline Luminosity system

Comments On The Overall System

EXECUTIVE SUMMARY

This executive summary consists of the main observations and main recommendations which are described in more detail in the rest of this document. An impressive and functional luminosity system has already been implemented thanks to the very detailed work and dedicated effort of few people. However, most of the Run 2 front end electronics have not been implemented yet and so the system is currently using Run 1 electronics. This system is currently being used in the experiment to determine luminosity online (real-time feedback to Beams Division and D0 Control Room) and for offline analyses. Below we divide the system up into four parts: Hardware, Readout/Simulation, Online, and Offline and give a summary of our observations and recommendations.

Observations Hardware:

Current system functionality is based on Run 2 scintillator counters, phototubes and preamplifiers combined with a partial reconstruction of the Run 1 electronics until the Run 2 electronics comes into operation. The LM counter performance is not monitored.

Recommendations Hardware:

1. We agree with the luminosity hardware group that additional manpower will be needed to bring the full Run 2 electronics past the prototyping stage into production and operation.
2. The highest priority should be given to finding an engineer to work with the Brown group on commissioning the electronics system and familiarizing him/herself with the system. This will provide needed support and form the basis of a future detailed engineering evaluation of the hardware system for DØ management. Additional collaborator and technical support should be identified to work with the hardware group on the following additional tasks:
    a. Test stand design, location and setup and
    b. Calibration and monitoring of system performance
3. A small committee should be appointed to monitor the progress of the Run 2 electronics and report back to the spokespeople.

Observations Readout/Simulation:

The measured luminosity currently has a 10% error assigned to cover uncertainty in efficiency and acceptance of counters, which is estimated based on Run I numbers.

Recommendations Readout/Simulation:

1. Additional collaborator support should be charged with getting the luminosity detector readout into the data path, integrating that data into offline reconstruction programs and producing a parallel Monte Carlo effort for data/simulation

comparison to enable efficiency and acceptance calculation to reduce systematic uncertainties in the final calculation.
2. This effort must be in coordination with the hardware group and feedback to that group doing comparisons of data from Run 1 and Run 2 electronics paths will facilitate the understanding of both systems.

Observations Online:

1. DØ has an extremely reliable, well designed, battle proven online luminosity DAQ and book keeping system
2. This system delivers a luminosity for every data set which is recorded under appropriate criteria
3. Online system in current form will be good for the remainder of Run II (except for small upgrades, maintenance and monitoring) and its resource needs are supported by the online group although web based report applications should be moved offline wherever possible. The online system is used to monitor many DAQ and trigger functions and is often the first responder for problems.
4. Changes to the system would be necessary to perform the current level of cross checks if online streaming of physics triggers is employed.

Recommendations Online:

1. The system is working very reliably, but is complicated. The documentation describing the system and all its features needs to be produced.
2. While the online luminosity system has been extremely useful in finding problems with the DAQ and triggering systems,
    a. Any additional online consistency checks needed between subsystems should be added to a global online administrative process and not the online luminosity system.
    b. Checking results and finding the source of inconsistencies should be the responsibility of the subsystem and not the luminosity group.
3. Additional manpower needs to be identified for changes to the system necessary to perform the current level of cross checks if online streaming of physics triggers is employed.
4. Web based luminosity reports should be moved to the offline system employing a dynamic web application utilizing the offline database (or equivalent offline system).
5. New manpower needs to be identified to support, upgrade and monitor the luminosity DAQ system.

Observations Offline:

1. Luminosity information is available for offline analysis using information in flat files produced from the online luminosity system.
2. Tools exist to access the information for analysis (“good” lum or “bad”)
3. Analysis Tools group is currently writing tools to handle data access with streaming and handle the complication of lost files.

Recommendations Offline:

1. The offline luminosity information should be stored in a database (scalability, satisfy all offline needs, etc.).
2. A design exists and it should be implemented as soon as possible. Currently database expertise from CD is available. Should take advantage of this now.
3. This database would contain more information than the current flat files and would serve all offline luminosity needs (from analysis to weekly and monthly reports, feed WEB pages, etc.)
4. To accomplish item 1) physicist help is essential and needed now. One physicist dedicated to offline luminosity should start working full time now, learning the database implementation and developing applications.

Observations Overall System:

1. Current luminosity system works very well and provides a luminosity number to the analyses adequate for current needs of the experiment including some book keeping, constantly recovering previously ill-defined luminosity blocks using new information)
2. The luminosity returned for analysis has inherent uncertainties which are covered by a 10% error, but which are undetermined

Recommendations Overall System:

1. An immediate influx of people is needed to get the offline system to be more robust and to make calculations of luminosity to the precision required for publishing cross sections from Run II. It requires a long time commitment by individuals and/or groups.
2. Documentation of the whole system needs to be produced

THE HARDWARE OF THE LUMINOSITY SYSTEM

Findings (Hardware)

The luminosity monitor hardware consists of 24 scintillation counters (12 on each side, Bicron 408, about 6" long in a trapezoidal shape). Each counter is read out by one mesh dynode phototube (Hamamatsu, with quartz window). Six scintillators are housed in an enclosure with a single 6-channel preamplifier board. The housing is roughly airtight with gas fittings to allow dry air to flow through the enclosure. Signals are brought to the counting room on high quality "RG8-like" cable. Presently the signals enter a Run I era "patch panel" that equalizes timing to the 100 ps level. From there the signal is passively fanned out to two places. The first is the Run I style electronics that is presently being used to produce the luminosity (NIM summers, "x10" amplifier, AC coupling, constant fraction discriminator, "fast z" module producing the z vertex signal). The second path is the one that is being used by the new VME based electronics (PowerPC, SBC, TDC cards, VTX cards, MFC cards) which is in the early stages of commissioning. The new electronics is designed around a Muon readout crate with custom TDC and VTX cards. Both systems can run simultaneously, thereby allowing the new electronics to be commissioned while still using the old electronics for luminosity. The space occupied by the system (both old and new simultaneously) is about 1.7 racks - the overflow is "borrowed" space but there does not seem to be any pressure to return it yet.

The bench measured timing resolution for the counters is at the specification level of 200ps. The non-uniformity in response of the counter/tube system has been measured on the bench -- it peaks at the location of the tube and falls off by a factor of ~10 far from the tube. The tubes sit in a 1T field (~axial with the tube) and one measures an ~30 change in gain between field on/off. This is as expected from the manufacturer specs.

The present status is as follows:

- Counters/Tubes: All counters/phototubes/preamps are in place and working. The cabling is in place (although perhaps too long).
- Electronics: The Run I style electronics is in place and working as expected. The new electronics ranges in status from commissioning (15 TDCs are complete and tested, of which 6 are required by the luminosity monitor, and they are presently being commissioned in the crate and integrated with the MFC), to near production (the VTX board has completed its second and perhaps final layout and will be sent for production soon). Turnaround for VTX board production and stuffing is ~1 month. One VTX card is required for the Lum system.
- Test Stands: Presently the debugging of electronics occurs in a luminosity VME crate located next to the old electronics system rack (at the bottom of the rack) in MCH1. Two other test stand areas have been considered. One is the muon test stand on the 2nd floor, but it is presently configured for muon use only (and not released to the luminosity group). The other is in the assembly building (in the area previously used

for AFE debugging), but it is completely undeveloped and in need of infrastructure before it could be used (rack, power, etc).

- Manpower: The tasks required to commission and operate the luminosity hardware have been estimated by the hardware group, and summarized in Table II of document entitled "Response to Questions from the Luminosity Review Committee". We reproduce that table here (assuming that completion will take 9 months from now):

Task	**FTE**	**Duration**	**Identified**
TDC/DAQ debugging	1 phys	36w	0.7 Casey
VTX board hardware	1 phys	36w	0.05 Casey + 0.1 FPD
VTX FPGA programming	1 EE	12w	0.05 Casey
Calib. Infrastructure	1 phys	26w	1.0 Bodyagin
Calib. Algorithms	1 phys	13w	none
Operations	1 phys	36w	0.1 Casey
Steady State Ops	0.75 phys	forever	0.1 Casey

Thus we conclude there is a present need for 4 FTE physicists for the next 9 months and 1 FTE EE for 3 months. The present identified manpower is 1.6 FTE physicists and no EE.

While not as urgent as the basic and immediate functionality of the system, there are a number of areas of concern mentioned in presentations which the committee agrees need to be addressed with additional manpower:

Helium Contamination - The level of Helium in the region of the luminosity counters is higher than anticipated (10-20ppm vs. natural level of 5ppm). This is an issue regarding the operation and lifetime of the phototubes (Rich Partridge notes that CDF is/has replaced tubes because of Helium contamination – regarding this statement : the CDF conclusion about this issue was that the malfunctioning of PMTs was compatible with He contamination but the contamination happened BEFORE the PMTs were installed in the collision hall. Possible reason could be contaminated isobutane. The CDF PMTs operate in a gas tight enclosure in atmosphere of isobutane (Cherenkov radiator) so the He contamination of the collision hall is not relevant in this case.). Although there are no signs of feedback in the PMTs, it is not known how severe the problem is. For example, the quartz window (for improved radiation damage) used on the tube (rather than the standard borosilicate glass) is significantly more porous to helium. Purge lines were installed in the boxes to mitigate the helium contamination problem, and the effectiveness will be able to be judged after a few Tevatron quenches.

Radiation Damage - The radiation dose is estimated to range from 25krad/fb-1 to 300krad/fb-1 for the scintillator depending on the distance from the beam. It is unclear what the effect will be, but the group estimates that the scintillator will not have to be replaced before RunIIb. The radiation damage to the PMT is not expected to be an issue. The dose in the region of the preamps in the enclosures is estimated to be ~15krad/fb-1. There has been no study of radiation damage to the preamp board (consisting of

commercial op amps and voltage regulators mainly). If the damage causes baseline shifts, then the TDC boards can compensate for the effect.

Spares - The group provided an inventory of spares and it would appear that for the systems for which they are directly responsible they have 25%-150% spares. There are a number of systems which are maintained by the Muon, Controls, and L3 groups for which they did not have definite numbers of spares.

Reliability - The preamps that are located in the luminosity enclosures are in an area that is difficult to reach (requires opening the detector). There is reason to believe this is acceptable because: (a) they have been running for 1.5 years with no failures; (b) a single channel failure can be tolerated; (c) if the preamps die, the fallback could be to remove them entirely and just carry the signals in on the cable - the noise is low enough to do that. The cables need to be checked every time there is an access to verify that there has been no excessive strain put on the cables and the strain relief. This is done routinely after every relevant access.

Comments (Hardware)

There are a number of areas where the committee had detailed comments which are discussed below. Overall the Committee was impressed by the presentations, and clearly heard the message that there was a significant lack of manpower to complete this project in a timely manner. It was quite concerned about the timescale outlined for the completion of the hardware, and felt that there was a critical need to monitor the stability of the luminosity monitor. The principal recommendation is that D0 management immediately identify a suitable electrical engineer to work with the Brown group in commissioning the electronics. This will serve two purposes: the first is to supply much needed help, and the second is to supply an independent detailed evaluation of the electronics (will the electronics function as designed? Will it be completed by the Fall of 2003? Etc) to D0 management.

The above recommendation is a reflection of the serious understaffing of the hardware effort. The group has estimated the need for 4 FTE physicists and 1 FTE EE for at least the next six months. They presently have identified about 1.7 FTE physicists and no EE. If this situation continues it will surely delay the project beyond the already late date of Fall 2003. It is a serious problem that needs to be rectified.

There was concern that the commissioning time estimates have large uncertainties. While the TDC boards were tested at Brown, the functions that could not be tested at Brown such as the SCL interfacing is not yet working properly, and is the focus of much of the present activity. To date there are a number of areas which could have potential problems because they have not yet been tested, e.g. VME registers have not been read, the ADC/TDC performance on the board has not been tested, the EPICs database is not setup for the new crate yet. Besides the firmware issues described above, there may be hardware issues lurking related to board layout, terminations, cross talk, etc. The committee was not in a position to evaluate the potential future problems, but was

concerned enough that they felt having a qualified engineer working with the group as soon as possible was essential to understanding the robustness of the system.

The VTX board is not yet in hand, so there remain concerns about the schedule for its production. In addition it was not clear to the committee how this board could be functionally tested until the TDC debugging was complete. At least beyond simple tests that verify that the firmware can be downloaded. There is also an obvious manpower shortage to test and commission this board that needs to be addressed.

The committee was pleased to note that the group has planned to allow the simultaneous use of the old system which is providing luminosity data together with the commissioning of the new electronics. The concern, however, was that the infrastructure and location of the test stands was inadequate. Presently the system is being commissioned and debugged in a rack in MCH 1. That environment is unsuited to extended debugging sessions, and perhaps inaccessible during running? Furthermore the instrumentation would appear to be inadequate as well as having the crate under test at the bottom of the rack - a very inconvenient location. Other locations had been explored such as the muon test stand on the 2nd floor and a location in the assembly building (previously occupied by the AFE testing). In the first case the muon test stand is still in use by the muon group and cannot be modified for luminosity use, in the second case the space needs significant preparation work - racks, power etc. We encourage the luminosity group to identify more suitable space, and we urge management to supply the necessary technical help to supply the requisite infrastructure. The number of spares are probably adequate (including those parts which require lifetime buys because of impending obsolescence), although it would be helpful to have a clearer understanding from the other groups that must support other parts of the luminosity system - such as the muon, L3, and controls group what their spare and maintenance policy is. This would avoid any potential problems.

One may ask whether there is a need for new electronics given that we seem to be getting luminosity measurements based on the old electronics? The luminosity group convincingly argues that the old electronics cannot be resuscitated to full functionality (or at least not without a major effort). Given that the new electronics has been in process for quite some time, it does, however, raise the issue of whether the design should be re-evaluated in light of electronics developments and design practices since the design was completed - in other words would a new improved design allow one to have the same functionality in a robust and easier to implement system? The committee did not feel qualified enough to proceed any further with that thought, but it remains one that future committees may want to consider once an engineering evaluation of the new electronics is in hand.

<u>Recommendations (Hardware)</u>

1. We agree with the luminosity hardware group that additional manpower will be needed to bring the full Run 2 electronics past the prototyping stage into production and operation.
2. The highest priority should be given to finding an engineer to work with the Brown group on commissioning the electronics system and familiarizing him/herself with the system. This will provide needed support and form the basis of a future detailed engineering evaluation of the hardware system for DØ management. Additional collaborator and technical support should be identified to work with the hardware group on the following additional tasks:
    a. Test stand design, location and setup and
    b. Calibration and monitoring of system performance
3. A small committee should be appointed to monitor the progress of the Run 2 electronics and report back to the spokespeople.
4. The group should identify a suitable location for electronics testing and debugging. They should be supplied with the necessary technical support to build up the appropriate infrastructure.
5. Additionally, attention is needed in the areas of detector monitoring and stability.

Comments (Readout/Simulation)

Because of a lack of manpower, there has been little or no work on the Monte Carlo work required to measure efficiencies and acceptances. This is the main reason why the error on any quoted luminosity is of order 10%. Effort in this area is needed immediately.

Recommendations (Readout/Simulation)

An immediate effort is required to determine efficiencies and acceptances in order to improve the luminosity measurement for D0 to meet it's design goals for precision measurements in Run 2. This effort must start very soon:

1. Work with the hardware group to add luminosity system measurements to the event record
2. Add to the offline reconstruction programs to analyze the luminosity system output
3. Write a Monte Carlo simulation program for comparison with data.
4. Feedback from this effort should additionally help to identify problems with the hardware and electronics.

THE ONLINE LUMINOSITY SYSTEM.

Findings (Online)

The online part of the system was described in a talk by Michael Begel. To measure the luminosity, the coincidences between the luminosity counters on the south and north side of the experiment are used as input. Depending on the time of the coincidence fast logic decides whether the coincidence is caused by "proton halo", anti-proton halo" or due to a real interaction i.e. "luminosity". The coincidence information is

accumulated in the per bunch scalers within the Level Trigger framework. Currently this is the only information being read out and stored from the luminosity counters. No further data are read out from the luminosity counters at his time (see hardware section). The luminosity data acquisition system acquires the above mentioned coincidences, by reading out the per bunch scalers associated with each Level trigger bit and for each of the 159 ticks of the accelerator. This way the luminosity and live fraction of each trigger for each bunch in the Tevatron is read out and stored by the datalogger in files on the online system. The stored data consists of so called Luminosity Blocks. A Luminosity Block typically contains the luminosity information for 1 minute of running time. Each is identified by a Luminosity Block Number (LBN) which is unique for all of Run II. More information on the online system can be obtained from Michael Begel's talk. The tasks performed by this system can crudely be summarized as follows:

- Read out the per bunch scalers
- Based on this information calculate the instantaneous luminosity
- Report the instantaneous luminosity back to ACNET for use by the Beams Division
- Store the Luminosity Block information in files on the online system
- Through examines enable monitoring of per bunch intensities and halo
- The system also performs a series of consistency checks for each LBN by looking for missing crates in the detector readout, checking that rates out of L1 and L2 are consistent with input and output rates of L2 and L3 etc.

The system has been in operation since July of 2000 (start of Run II was March 2001), it contains GUI's to display and control the HV of the luminosity counters, there are several levels of GUI's to monitor the performance and operation of the luminosity DAQ system, there are recovery procedures to either recover information or repair bunch scaler information, a series of consistency and validation checks are performed regularly to ensure that the data taken is of good quality and that actual data taking runs are normalizable. The system has grown into a mature, very reliable separate DAQ system for readout and storage of luminosity information. A lot of effort has been put in to recover information and on the online side this system provides luminosity for any data set that the experiment has recorded and will record. In summary it is an excellent system that has proven to be extremely reliable under the varying operating conditions from the accelerator and from the experiment over the last two years. The system in its current form will serve the needs of the experiment for the remainder of Run II. The author of the system foresees small upgrades in the future like the format used for the communication path with ACNET. No other major upgrades are foreseen or planned. However the system will have to be maintained and monitored and this will require some manpower.

Because of the need to create normalizable data runs the "luminosity group" also performs consistency checks of the standard DØ DAQ system (see above) to make sure that the data written to tape are internally consistent. Examples are checks for loss of data, etc. In the past and currently the luminosity group therefore performs an important service to the DAQ group in terms of checking the data consistency. The luminosity group is often the first to discover problems in the DAQ and triggering systems.

Because the system calculates the delivered and recorded luminosity, as well as sources of dead time it is also used to produce weekly and monthly operations reports. These reports reflect which fractions of the delivered luminosity was used by DØ and therefore defines the data taking efficiency of DØ. These reports are presented at weekly experimenters meetings, at accelerator division meetings and are available to the collaboration on the lum ID WEB pages..

Although originally it was planned to store the Luminosity Blocks and other information in a database on the online system, the database never materialized. So the data are stored in flat files on the online system. This system seems to be working very well and the committee sees no reason to suggest any changes here except for the replacement of the online web page support (flat html files) with the offline database and applications (dynamic html).

Observations (Online):

1. The online luminosity system is working with a very high degree of reliability, it performs all the tasks required by such a system (collecting the data, monitoring the data and verifying the quality of the data, making the data available online while running and for offline reports) and the authors should be commended for designing and implementing such a system.
2. Online system in current form will be good for the remainder of Run II (except for small upgrades) and its resource needs are supported by the online group although web based report applications should be moved offline wherever possible. The online system is used to monitor many DAQ and trigger functions and is often the first responder for problems.
3. Changes to the system would be necessary to perform the current level of cross checks if online streaming of physics triggers is employed.
4. The current system is not equipped to track Luminosity Blocks in the case where triggers are written to multiple streams when some streams are inadequately populated on a per block level. Note that the Data Logger has been upgraded to provide information about rare streams. The online luminosity system needs to be upgraded to use that information to discern a rare stream from lost data.
5. One foreseen upgrade is the communication with ACNET
6. The online luminosity system has so far acted like an “administrative process”, performing consistency checks between online subsystems (e.g. comparing L2 accepts to L3 inputs). The luminosity ID group has done an excellent job identifying such discrepancies and inconsistencies, and often pursuing the solutions when the online subsystem groups have been unwilling to do so.

Recommendations (Online):

1. The system is working very reliably, but is complicated. The documentation describing the system and all its features needs to be produced.

2. If online streaming of physics data is expected, provide manpower to upgrade the system to address the accounting problems (e.g. rare streams) that might arise.

3. Upgrade the ACNET communications, as proposed by the authors

4. While the online luminosity system has been extremely useful in finding problems with the DAQ and triggering systems,
    a. Any additional online consistency checks needed between subsystems should be added to a global online administrative process and not the online luminosity system.
    b. Checking results and finding the source of inconsistencies should be the responsibility of the subsystem and not the luminosity group.
5. Additional manpower needs to be identified for changes to the system necessary to perform the current level of cross checks if online streaming of physics triggers is employed.
6. Web based luminosity reports should be moved to the offline system employing a dynamic web application utilizing the offline database (or equivalent offline system).

7. New manpower needs to be identified to support, upgrade and monitor the luminosity DAQ system. The ideal solution would be a university group, which would take on this responsibility. The total estimated manpower is 0.5FTE's.

THE OFFLINE LUMINOSITY SYSTEM

Present Situation (Offline)

Determination of luminosity requires three distinct steps.

- Determination of correspondence between data files and luminosity blocks
- Ensuring that the luminosity block list actually corresponds to the analyzed data (normalization)
- Determining the total integrated luminosity given the list of blocks

Since luminosity information is not stored in the data files themselves, a system is required to determine the correspondence of luminosity to the data files. Currently, information about a *luminosity block* (on average about a minute's worth of data taking) is saved in flat text files that are produced on the online system. There is one "stage2a" file for each LBN and it contains status and luminosity information for each trigger, results of offline reconstruction checks (performed by an automatic offline process that are fed back to online), and raw instantaneous luminosity information. Files for good LBN's are copied to the offline systems (d0mino and clued0). Files for bad luminosity blocks are not copied. All of the good stage2a files sum up to about 1.5 GB (gzipped). There are about 70,000 luminosity blocks in the Moriond 2003 sample.

The *lm_access* code handles access to this luminosity information but does not translate files to luminosity blocks. It is a stand-alone program that takes a list of luminosity blocks, does rudimentary quality cuts on the LBNs, and returns the integrated luminosity. It is accessible from CVS. There are scripts available to determine the luminosity blocks that correspond to a set of files. There is a new first version of *lm_access_pkg* that is a framework package for determining luminosity by file. It uses the SAM database to determine the luminosity blocks associated with a file (so this only works for files that have been stored in SAM). Right now, this process is very slow, but a significant speed up has been identified. *lm_access_pkg* is then a wrapper for *lm_access*. The package is in the regular DØ releases.

Normalization is the process to ensure that the calculated integrated luminosity for a data sample actually corresponds to the data analyzed. For example, if *lm_access* rejects a luminosity block, then the user must also reject the corresponding events in their analysis. Furthermore, *lm_access* must know the user's bad run list to not count luminosity for bad runs. There are other issues regarding checking for duplicate events and ensuring that the user is analyzing files reconstructed with an appropriate version of RECO. Currently, *lm_access* does some of these tasks for a fixed RECO version and assuming all events are written to a single stream or that all files from all streams are processed. The calculation of the integrated luminosity is then a matter of summing luminosity for the desired trigger set for the good luminosity blocks. The stage2a files contain the calculated integrated luminosity for each trigger. If a luminosity constant changes, then the stage2a files are regenerated by the online system and copied to offline. *lm_access* performs the luminosity calculation.

Future (Offline)

The current *lm_access* software cannot handle streaming. Difficulties with streaming include handling multiple files that correspond to the same luminosity blocks and the more challenging normalization checks. Given that much of *lm_access* will need to change to accommodate streaming, the luminosity group sees now as a good opportunity to incorporate other changes: using a database for luminosity information instead of the flat files and obtaining run quality data from the run quality database. Exclusive streaming means that data from a trigger may be present in more than one stream. In order for a luminosity calculation to be accurate, the user must be sure to analyze all of the streams that could have received events. If one necessary file for a particular stream is not analyzed, then there could be a biased loss (since data in different streams have different topologies). Note that the unrecoverable file loss rate from RECO and SAM are small (reco is < 1%, permanent SAM losses are nonexistent). But a few files for very recent runs may not be present in SAM or undeliverable due to reco failures (this recoverable loss rate is about 2%). The reco farms do eventually go back and recover this lost data. If this biased loss were not noticed by the user, the luminosity would be overestimated and data based efficiency studies could be skewed. If one knew about such a biased loss, one would determine the luminosity blocks in the missing file and remove events in those blocks from the successfully analyzed data sample. Now the biased loss

has become simply a decrease in seen luminosity. The Analysis Tools Group is writing a tool to discover such a biased loss.

The flat text stage2a files do not scale well with increasing luminosity (there are problems of having too many files in a directory) and they are not self-describing. A luminosity database in DØ's Oracle system is being developed by the luminosity group and Jeremy Simmons, a DØ database consultant. Jeremy will write first versions of applications that use the database, but it is not clear who will make and maintain a production system.

The luminosity database is an extension of the stage2a files. Instead of the results of the calculation to determine integrated luminosity for each luminosity block (by trigger), the raw components are stored in the database, along will all necessary constants. Storing the raw scalar data allows for dynamic updating the luminosity when a scalar value or constant is changed (so whole repopulation of the database should never be necessary). Additional information is stored for obtaining instantaneous luminosities and trigger cross sections. Many status words are kept to enable cuts on luminosity block quality and reconstruction completeness (e.g. did RECO process all events in the file?).

There is no versioning information planned for the database, except for perhaps the update time and user.

The luminosity database will reside in DØ's mission-critical Oracle DB system. The system is automatically backed up regularly and has 24/7 support by Fermilab staff. The luminosity group currently has no personnel for developing the database and its necessary applications and no one for maintenance. They estimate they would need 2.5 FTE's for 6 months (not including Jeremy), with at least one FTE being a physicist.

<u>Recommendations (Offline)</u>

1. The offline luminosity information should be stored in a database for improved scalability, added functionality, and support from Fermilab staff.
    a. Jeremy Simmons is in the process of implementing the current database design. In the past, the database design effort has been incoherent, despite Jeremy's best efforts, due to lack of a primary physicist contact for design changes and absence of development on the client side (user applications).
    b. New physicists manpower must be identified to take responsibility for this project. Primary focus must be dedicated to immediate development of the needed applications as the driving force for design changes, which has been the missing element in previous implementations.
2. A design exists and it should be implemented as soon as possible. Currently database expertise from CD is available. Should take advantage of this now.
    a. New manpower, with Jeremy's assistance/advice, should directly address the following questions: How will data be entered into the database? Web pages? Automatic scripts? How will the run/daily/weekly/monthly reports

be produced? What tools (aside from *lm_access*) will be available for users to investigate their luminosity calculations?

b. While the stage2a files are a complete black box to users, the database need not be so. Users should be able to investigate their luminosity results with some tools to browse the luminosity database. The online system will continue to keep the text files that hold the master luminosity information. How will these be connected to the database?

c. This database would contain more information than the current flat files and would serve all offline luminosity needs (from analysis to weekly and monthly reports, feed WEB pages, etc.)

3. The lack of versioning in the current database design is somewhat worrisome. Users will ***not*** be able to reproduce previous luminosity results unless they cache database information locally (perhaps the new *lm_access* system will do this for them). While this idea will work, it is undesirable because of the lack of reliability of such information. Furthermore, without luminosity block versioning, there is little protection against data corruption (aside from the regular backups) by errant database updates. Users whose luminosity drastically changes for no apparent reason will have to talk to the current luminosity expert to find out what happened. But adding versioning information to the database appears to be very complicated, and will not be in the first implementation. While it would be more prudent to start off with a schema that mirrors the stage2a flat files, additional information and the raw components should be stored in the database to maximize help from Jeremy while he is assigned to the project.
4. Finally, the manpower issue is very disturbing. The presented estimate of 2.5 FTE's does seem highly conservative. One sure-fire approach would be to have one physicist dedicated >90% to offline luminosity, working full time with Jeremy Simmons learning the database implementation and developing applications. After 3 months, a more accurate assessment of the additional manpower required could be made. This person and any additional others would continue in the offline luminosity effort when Jeremy is no longer contracted to work on this project.

COMMENTS ON THE OVERALL SYSTEM

Despite the delays of the Run 2 electronics, the current luminosity hardware works well enough for now and provides a luminosity number to all analyses. This is achieved by book keeping and constantly updating the numbers. The luminosity returned though has inherent uncertainties which are believed to be covered by a 10% systematic error. This error has to be reduced so that cross sections can be measured accurately. To reduce this error, the efficiency AND acceptance of the luminosity counters has to be determined, which requires reading out the counters and simulation of the acceptance of the counters for the processes making up the cross section seen by the counters. This effort has been partially described in the hardware section already and the work on the hardware is all geared towards reading out the counters. To do the remaining work (simulation and inclusion of the counters in the regular data stream) more people are needed. This is another area where an immediate influx of effort is needed.

The question was raised during discussions whether the system efficiency could be determined in an alternate way, by using standard data sets. At least this might be possible for the relative efficiency over time. This may be possible and might be worthwhile pursuing, until the hardware and software are in place to measure the efficiency directly.

The need to access the offline luminosity information in remote institutions has become rather apparent. Collaborators are actively analyzing data outside of Fermilab and they now need access to luminosity, because it is an integral part of the analysis chain. The current format of storing the offline lum information (flat files) and the method for keeping them up to date are not easily transportable and scalable to an analysis system which is location independent. This is another argument for the need of a database to contain the offline luminosity information, which is accessible from outside Fermilab.

In an ideal world the lum ID group should determine and make available for analysis the lum corresponding to the data sets created by the online system. However once the data are shipped to the offline system they are manipulated further in a way which can alter the normalization for the luminosity for a data set (loss of data or up coming streaming are examples). In the current system the lum ID group has taken on the task of feeding this "normalization" information back and re-determining the lum for offline manipulated data sets. This is referred to as "reconstructed" luminosity. This seems awkward. As the complexity of the offline system increases, it may be worthwhile considering forming a "Cross Section" ID group. The task of this group is to make sure that whatever DØ does with its data, it is done in a way so that it is still possible to derive cross section for physical processes. A subgroup of this cross section group would be the lum ID group, whose responsibility it is to provide the online lum "only".

<u>Recommendations Overall System:</u>

1. An immediate influx of people is needed to get the system to be more robust and to make calculations of luminosity to the precision required for publishing cross sections from Run II. It requires a long time commitment by individuals and/or groups. The additional manpower is needed in the following areas (this is a summary of previously stated manpower needs in each subsection):
    a. To help the Brown Group commission components of the Run 2 luminosity system electronics to complete the Run 2 luminosity data path
    b. Methods for calibration and monitoring of system performance must be established
    c. Monte Carlo simulation of the luminosity system must be developed
    d. The offline reconstruction program should be modified to include information from the luminosity system for comparison to simulation (to measure the real acceptance of the detector) and measure efficiency
    e. The online luminosity DAQ system must be maintained, monitored and made robust for all expected Run 2 data taking scenarios
    f. An offline luminosity database should be developed in coordination with file delivery systems to enable precise, robust and optimized luminosity

calculations for appropriately derived data samples.

2. The formation of a “cross section” ID group should be considered, which contains the luminosity group as a subgroup.

3. Documentation of the whole system needs to be produced

Appendix

Charge to committee:
The committee should conduct a vertical review of D0's plans for luminosity determination for physics analysis. The review should cover all components of the system: hardware, electronics, data acquisition, storage and processing of luminosity information, and tools for physicist access. A draft report will be available February 14th and a final report March 1st.

The review should attempt to answer the following questions:

- Does the system, as conceived, meet DØ's needs for physics analysis? Does it support all the kinds of analysis we will carry out?
- Is the system robust (does it allow us to recover from loss of luminosity data, from loss of tapes, streams etc.)
- Are the present resources adequate to deliver the system in a timely manner? Are the appropriate priorities set?

Committee members:
Elizabeth Gallas, Marumi Kado, Greg Landsberg, Adam Lyon, Andrei Nomerotski, Paul Slattery, Mike Tuts and Harry Weerts (chair)

Meetings and presentations:
The committee and the lum ID group met for the first time on December 6, 2002. A this meeting Heidi Schellman gave an introduction to the luminosity system. There were two more meetings after this. One on January 23, 2003, where Brendan Casey and Michael Begel covered the hardware and online parts and one on February 7, 2003 where Marco Verzochhi , Heidi Schellman and Adam Lyon covered the offline side of luminosity. There was an additional "committee only" planning meeting on January 31, 2003.

All presentations and notes can be either found on the luminosity ID page or on the luminosity review Web page at :
http://d0seververver1.fnal.gov/projects/meetings/lumreview/lum_review.htm

As part of the review M. Kado prepared a document outlining the need for luminosity in each of the physics groups over the next year. This document can be found on the WEB page above or directly at:
http://d0server1.fnal.gov/projects/meetings/lumreview/physics_view.pdf

The first summary of the findings of this review was presented at the February 2002 collaboration meeting. This talk can be found on the above WEB page or directly from:
http://d0server1.fnal.gov/projects/meetings/lumreview/lum_review_feb_2003.pdf